NASA-TM-112063

The Role of Graphics Workstations in Supercomputing

RNR-87-003

T. A. Lasinski
Manager, Workstation Subsystem
NAS Systems Division
NASA Ames Research Center
Moffett Field, CA 94035

Abstract

The use of graphics workstations in NASA's Numerical Aerodynamic Simulation Program is discussed. The NAS workstation environment consists of over 30 Silicon Graphics Inc. graphics workstations networked to a Cray 2 and other computers via both HYPERchannel and Ethernet. Present application software is described with particular emphasis on distributed graphics between these workstations and the Cray 2. A video tape illustrating a typical distributed application is available. Present and desired data rates are discussed. An attempt is made to extrapolate this experience over the next few years.

I. Overview of the NAS System.

NASA established the Numerical Aerodynamics Simulation (NAS) Program at its Ames Research Center with three principal goals in mind:

1. provide national computational capability to insure continuing leadership in computational fluid dynamics (CFD);

2. act as an agency pathfinder in the integration and use of large-scale computer systems through the use of state-of-the-art computer hardware and software technologies;

3. provide a strong research tool for NASA's Office of Aeronautics and Space Technology.

The NAS System is described in greater detail in ref. 1. Here only those features of NAS relevant to graphics workstations and supercomputing are discussed. Figure 1 shows all the computers associated the NAS System and how they are networked together. The Cray 2 has a 4 nsec. clock and 256 million 64-bit words (2000 million bytes!) of main memory. The four CPU's can obtain a peak speed approaching 1000 MFLOPS for vector calculations. Connected to the Cray 2 are many Silicon Graphics "IRIS" workstations. Both HYPERchannel and Ethernet are used. An ever increasing number of remote sites are connected to NAS over terrestrial lines ranging from 56 to 224 Kbps. For the last two years, NASA's Langley (LaRC) and Lewis (LeRC) Research Centers have had IRIS workstations connected to the NAS System.

II. Workstation Requirements for the NAS Program.

The NAS graphics workstations were chosen through competitive procurement in the spring of 1984. The principal requirements were as follows:

1. stand-alone CPU with VAX (780) performance
2. real-time, dynamic graphics capability
3. large disk space
4. good, direct communications to the Cray
5. capability to make video animations
6. tools to made "journal quality" output

These requirements are no doubt common to all potential workstation users. The second one was of particular interest in that it distinguished this workstation from the vast number of commercial workstations available then. While bit-map graphics was common place among workstations, the requirement called for the capability to manipulate graphics objects in near real-time. Today there are still only a small number of vendors who can deliver this capability in a workstation environment.

NAS users are for the most part working in an area referred to as computational fluid dyanmics (CFD). This discipline attempts to solve the differential equations of fluid flow numerically to understand flows over realistic three dimensional objects. Historically fluid dynamicists have studied flows with photographs using a wide array of techniques to visually enhance salient features of the flow. The use of computer graphics is a particularly natural step for the CFD scientist. The supercomputer provides him with a "wind tunnel" to perform his numerical simulations; workstations with dynamic graphics capability provide him with numerical flow visualization. Such applications are well described by Buning and Steger, ref. 2.

III. Description of Present Workstation Hardware.

A block diagram of the NAS workstation is shown in Fig. 2. Silicon Graphics, Inc. (SGI) refers to it as the IRIS 2500 Turbo workstation. "IRIS" stands for "Integrated Raster Imaging System". The Motorola 68020 chip and the Weitek chip set give the workstation floating point capability greater than that of a VAX 780 with floating point assist.

The unique feature of the IRIS is the Geometry Engine (TM). It consists of 12 VLSI chips which perform the floating point calculations needed to transform, project and clip geometrical data for display on a CRT. The IRIS performs transformations at a speed which requires Geometry Engine performance in excess of 10 MFLOPS. Some rendering of the graphics, including z-buffering (hidden surface) and smooth shading, is also done in hardware. The 32 bit planes of 1 Kb by 1 Kb resolution display memory can be configured in several ways. For animation two buffers of 12 planes each are used in a "double-buffer" mode. Full RGB color is obtained by using a single 24 bit plane buffer. Hidden surface removal is provided by using 16 planes for a z-buffer and 16 planes for color.

In addition to the usual compliment of RS232 and Ethernet ports, the IRIS also supports a HYPERchannel interface. This is simply a multibus board made by IKON (Seattle) which emulates a DR11W, an interface which permits connection to a Network Systems Corp. (NSC) A400 HYPERchannel adapter. This interface permits the IRIS to have direct communications to the supercomputer. A four trunk A400 cost about $80K andcan support four IRIS's. The connect cost per IRIS is thus $22K, including the IKON interface.

IV. The Network Environment.

There are three principal networks associated with the NAS system. The local area network based on Ethernet links all computers with the exception of the supercomputer. Cray will no doubt make Ethernet access available real soon now. The lack of an Ethernet interface to the Cray and the requirement for direct access of all computers to the supercomputer leads naturally to the use of NSC's HYPERchannel product. Both these LAN's connect computers located within the NAS Facility as well as to computers located in other buildings located at Ames. These campus wide connections are made via fiber optics, repeaters, and bridge connections for both Ethernet and HYPERchannel.

The long-haul or wide area network, NASnet, provides Ethernet access for over 70% of NAS users. This network connects the local NAS Ethernet to remote site Ethernets via Vitalink TransLAN communication bridges. The necessary terrestrial communication links are provided by NASA's newly implemented Program Support Communication Network (PSCN). This long-haul network has been in a prototype mode for over two years now. In particular, IRIS workstations located at Langley (LaRC) and Lewis (LeRC) Research Centers have Ethernet access to the NAS network.

All the computers at NAS run tcp/ip communication protocols. In addition Berkeley style networking commands are supported on all systems. This greatly facilitates the addition of new computers to the network. The use of internet protocols provides remote users transparent access to the Cray 2 via Ethernet-HYPERchannel gateways.

V. The Cray 2 as a Graphics Coprocessor.

In the foregoing sections, the tools for large scale scientific computing have been described. These are supercomputer, graphics workstation and the communications linking them. An application program known as RIP --- remote interactive particle tracer --- was written at Ames to use these tools to increase the productivity of CFD users. The UNIX and communications system issues associated with RIP are discussed by Choi and Levit, ref.3; two specific applications, RIP and PLOT3D, are described by Rogers et al. in ref.4.

Figure 3 illustrates the RIP application. It involves two processes, one on the IRIS and one on the Cray 2. These processes communicate over Ethernet or HYPERchannel using tcp/ip protocols. The process on the IRIS controls a graphics data base for the object under study such as a space shuttle or new fighter design. The IRIS can rotate and zoom through this data base in near real time, completely independent of the Cray. The IRIS also provides the principal interface to the user in this application. The process on the Cray 2 controls the solution data base for the object under study. This data base is the result of solving partial differential equations which describe fluid flow. It is typically 50MB in size and required 1 to 20 hours of Cray time to calculate.

The CFD scientist uses RIP by indicating with the mouse a point on or near the object where he would like to release a test particle. The solution data base has the information to show him how the test particle flows past the object. When he clicks the mouse, the location of the point is sent to the Cray. It in turn interpolates through the solution data base to figure out where the particle will flow. This technique is referred to as particle path tracing. The Cray 2 is very useful for finding particle paths. The search procedure is CPU intensive, and the data base is 50 MB or larger. While particle tracing can be carried out on the workstation, the process would hardly be considered interactive. In roughly 1/10 of a second, the Cray returns to the IRIS a series of perhaps 400 short vectors which geometrically define the particle trajectory.

In a matter of minutes the CFD user can define and build up a visualization of the flow field over the object of interest. At any point he can rotate the object and the traces to study the flow from different orientations. In this way he can correlate what happens in one area of flow with another. RIP provides him an interactive tool with which to visualize, to explore the results of his flow field calculation. Since RIP sends display list information, not image data, to the IRIS, all viewing and manipulation of geometry and traces is done independently of the supercomputer.

A video tape illustrating how RIP works is available upon written request.

VI. Data Rates

RIP can be effectively used by remote users since it does not send hugh amounts of data to the user's workstation. One trace is typically made of about 400 vectors. The tip of each vector is defined by 12 bytes (x,y,z). Another 3 to 4 bytes are needed for control and color information. At 15 or so bytes per vector, one trace amounts to about 6 kilobytes (KB) of data to be transmitted. This takes about 1 second at 56 Kb/s and 1/30 second at 1.5 Mb/s. For most users an interactive response corresponds to 1/4 second or faster.

The problem with particle traces is that they only reveal a fraction of information available in the simulation. In two dimensions CFD scientists often use contour plots of density or Mach number to gain a more detailed insight into the nature of the flow field. Such plots can be thought of as being made of many traces, typically on the order of 40 to 100 traces. In three dimensions users now stack 2D contour plots. In time more sophisticated visualization techniques will be developed. For now one can estimate that the equivalent of about 1000 traces will be involved. Roughly speaking then CFD users would like to view their data with single graphics frames corresponding to about one megabyte (1 MB) of graphics data. Such views would require 2.5 minutes to transmit at 56Kb/s or about 5 seconds at 1.5 Mb/s. These image densities (1 MB) are much more typical of what the CFD scientist would like to study interactively. In this regard RIP is a compromise and a fortunate situation. The CFD scientist can get some very useful information from the particle traces. This technique may very well be of little or no use to the visualization of other physical problems.

The visualization requirements discussed so far correspond to steady state flows where the flow field is independent of time. The are no end of unsteady flows which the CFD scientist would like to study. As a dramatic example of such a flow, consider the problem of store separation from a highly manuverable fighter. In many cases new store designs or new tactics lead to the store destroying a portion of the aircraft wing. The study of such phenomena is experimentally very expensive. The time dependent nature of this flow is clear. To visualize it would require hundreds of time steps, with each frame requiring 1 MB or so of data as discussed above. The motion of various components of this flow relative to one another is a critical feature of such a problem. This requires animation on the order of 10 frames per second. This sort of application thus requires data rates on the order of 10 MB/s.

As in other applications, the CFD user also wants the ability to rotate, zoom and pan through these images either on the fly or in a temporarily paused state. The work of Winkler et al. at Los Alamos, ref. 5, illustrates the power of animating CFD results at rates of 60 Mb/s. In this environment graphics images are stored on magnetic disk and transferred as quickly as possible to a frame buffer, providing animation of about 15 frames per second. This approach, however, does not permit the user to interact with the visualization. New external bus technology, such as that of the Ultra Corp., ref. 6., gives the promise of streaming data from the supercomputer itself to frame buffers at the rate of 100 MB/s. In this case interactive control can be gained with appropriate software on the supercomputer. This would constitute the world's most expensive graphics workstation!

VII. Concluding Remarks

An example of a distributed scientific graphics application known as RIP was discussed. This application uses a graphics workstation for the display and control of geometrical data and a supercomputer for CPU and memory intensive rendering operations. Workstation and supercomputer processes are closely coupled with standard networking software based on tcp/ip protocols. The use of a technique referred to as particle tracing keeps data transmission to a minimum. This in turn makes RIP a viable tool for remote graphics workstation users.

As discussed above it is evident that applications such as RIP can be quite successful in wide area networks running at T1. Although faster rates are always desirable, a truly robust T1 network would not only handle RIP-style applications, it would also permit sizable solution files to be shipped to the remote user's site for more exhaustive study. The word robust is very import. Scientific users want to move large (50 MB) files routinely to and from their home sites. The impression of many users with regard to today's existing 56 Kb/s networks is that they work fine for mail but are less than adequate for moving files.

While T1 networks may be suitable for the next few years, there are several factors pushing for greater capabilities. As graphics workstations become more common place, scientists will develop new techniques (2) to visualize their science. These techniques will require wide area rates in access of 1 MB/s. There already exist examples (5) of applications where local rates of over 60 Mb/s have been achieved. Graphics workstations in the next few years will be enhanced significantly in hardware capabilities. Multibus will be replaced by VME. Processing speed will go from the 1 MIPS regime to 10 MIPS or better. In stead of 100K vector transformations per second, there will be systems that transform 300K polygons per second including z-buffering (hidden surface). The advent of 4 Mb RAM chips will bring display memories of 256 bit planes. It is clear that future graphics workstations will have the fast buses and large memories needed to handle intensive supercomputer output. The question is whether future wide area networks will have the hardware and software capability to sustain communications between future supercomputers and graphics workstations.

References.

1. F.R.Bailey, "Status and Projection of the NAS Program", NASA TM 88339, July 1986.

2. P.Buning and J.Steger, "Graphics and Flow Visualization in Computational Fluid Dynamics", AIAA-85-1507-CP, AIAA 7th Computational Fluid Dynamics Conference, July 15-17, 1985.

3. D.Choi and C.Levit, "An Implementation of a Distributed Interactive Graphics System for a Supercomputer Environment", NASA Ames Research Center, Oct. 1986; to be published in the International Journal of Supercomputing Applications.

4. S.E.Rogers, P.G.Buning, and F.J.Merritt, "Distributed Interactive Graphics Applications", NASA Ames Research Center, 1986; to be published in the International Journal of Supercomputing Applications.

5. K-H.A.Winkler, J.W.Chalmers, S.W.Hodson, R.L.Judd, M.McGowen, D.E.Tolmie, "The Ultra-Speed Graphics Project of the Advanced Computing Initiative at Los Alamos National Laboratory", Los Alamos National Laboratory, 1986.

6. J.N.Perdue, Ultra Corporation, 2140 Bering Drive, San Jose, CA, 408-922-0100, 1986.

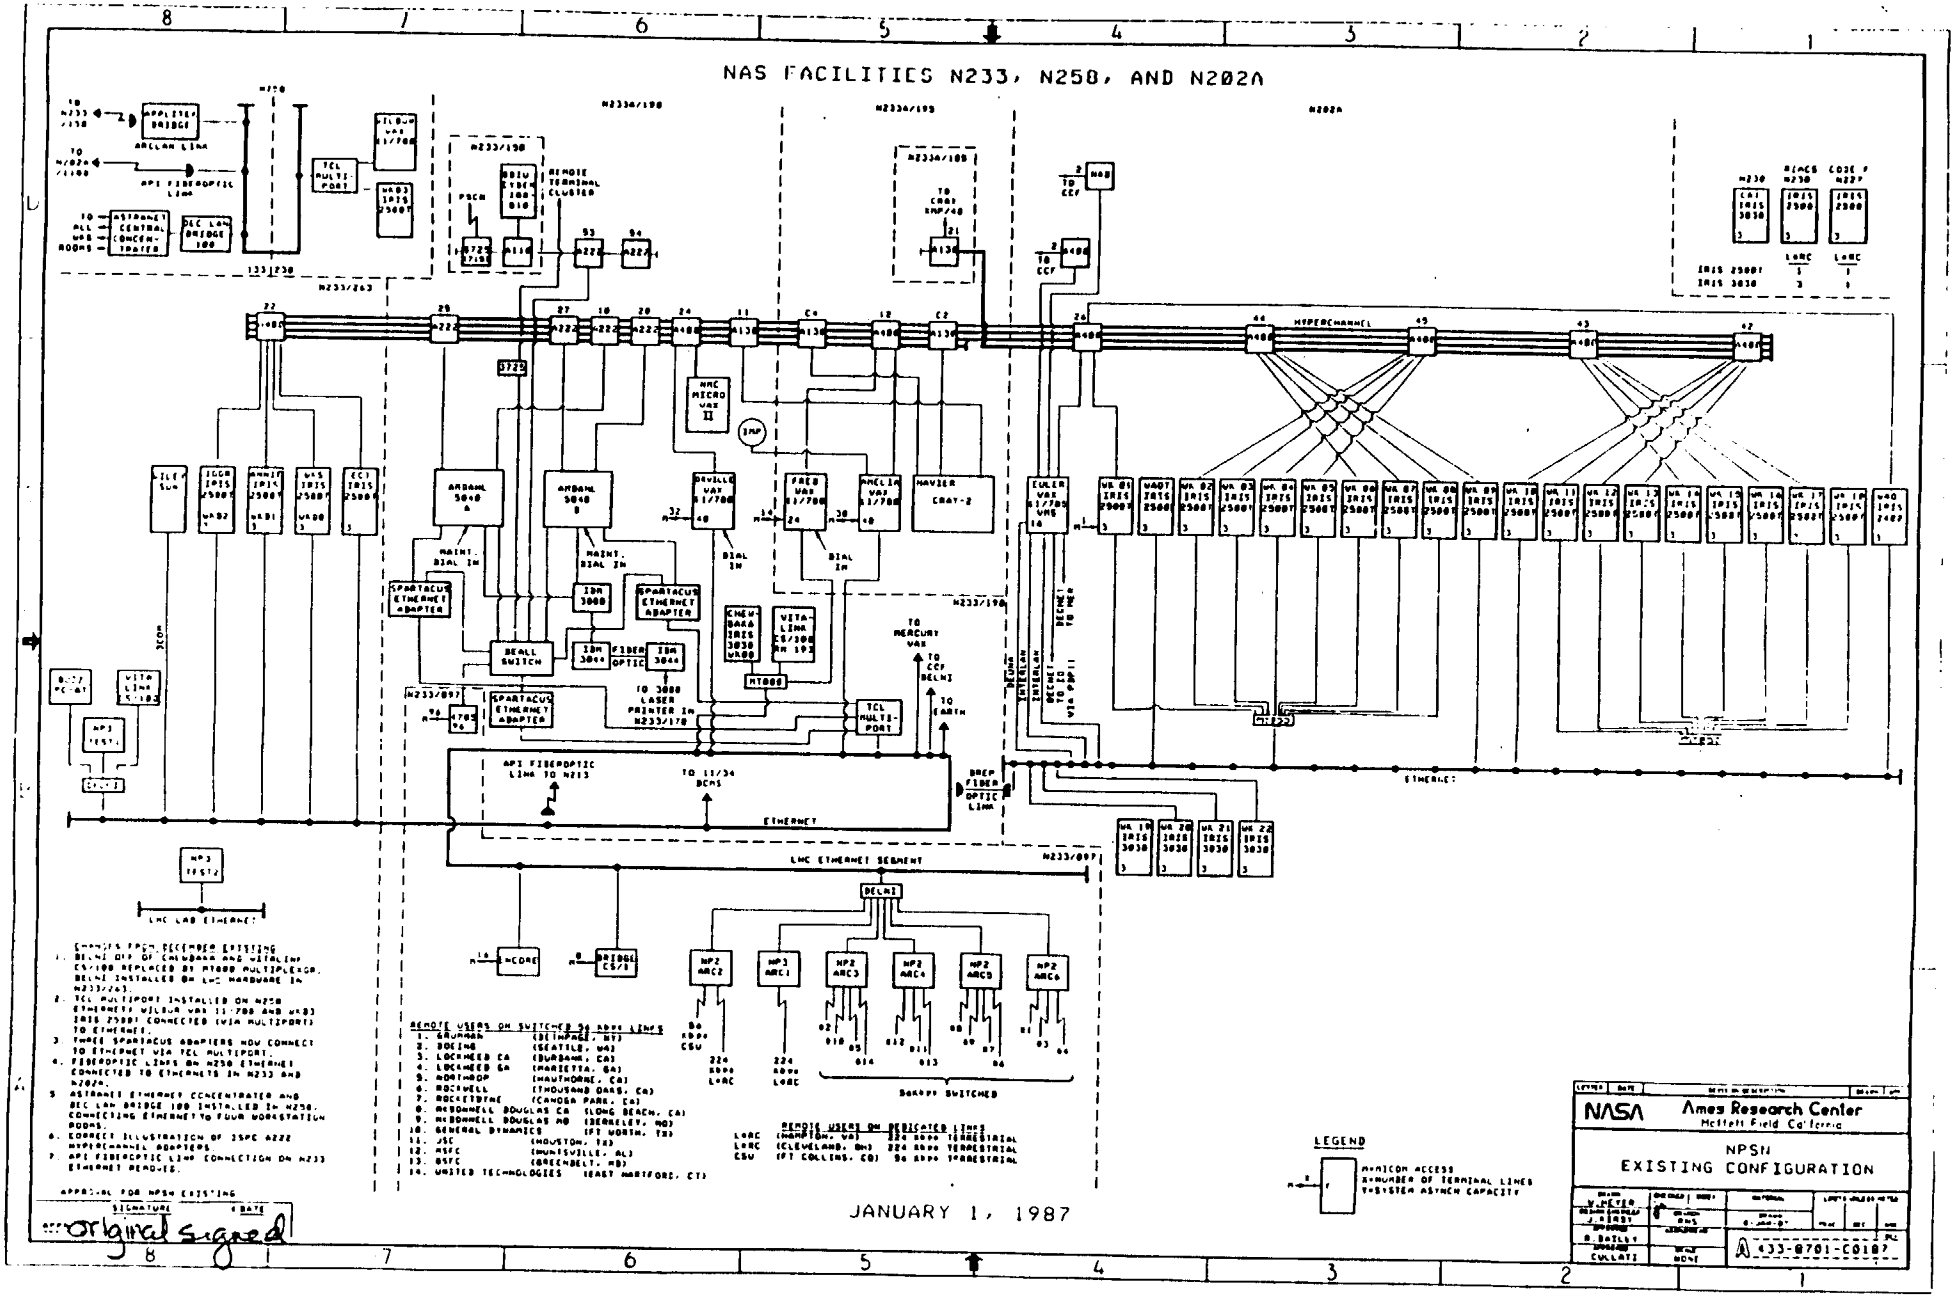

Fig. 1. Block diagram of NAS computer network

SILICON GRAPHICS (SGI) TURBO GRAPHICS WORKSTATION

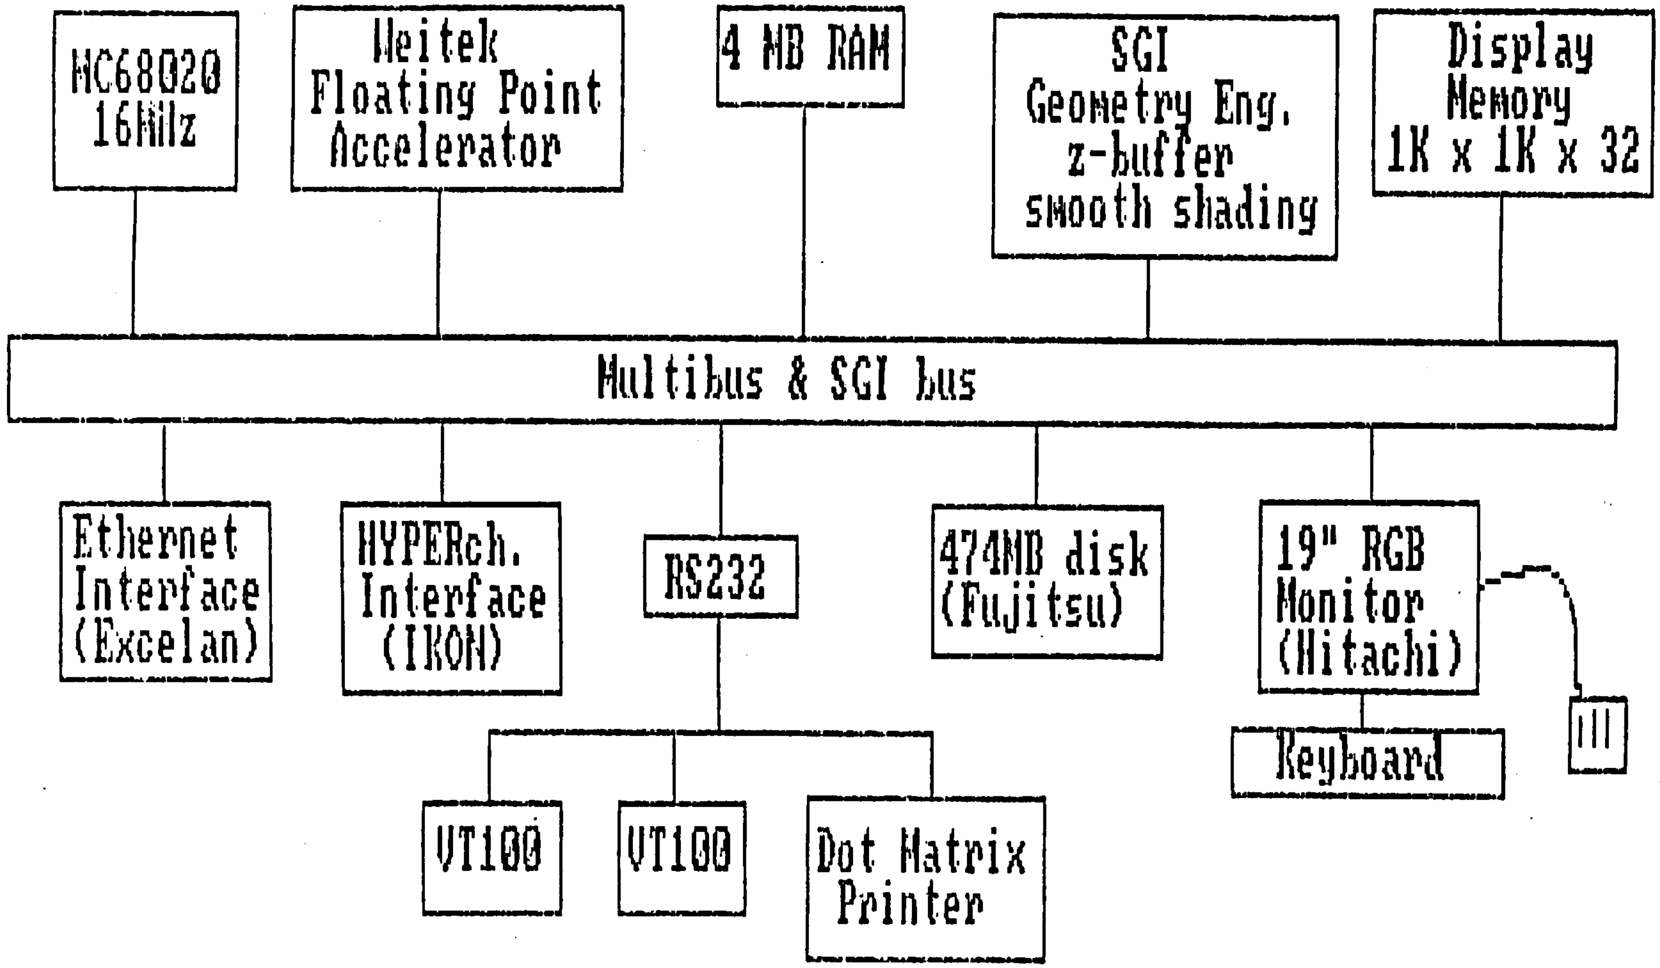

Fig. 2. Block diagram of NAS graphics workstation.

Lasinski - 3

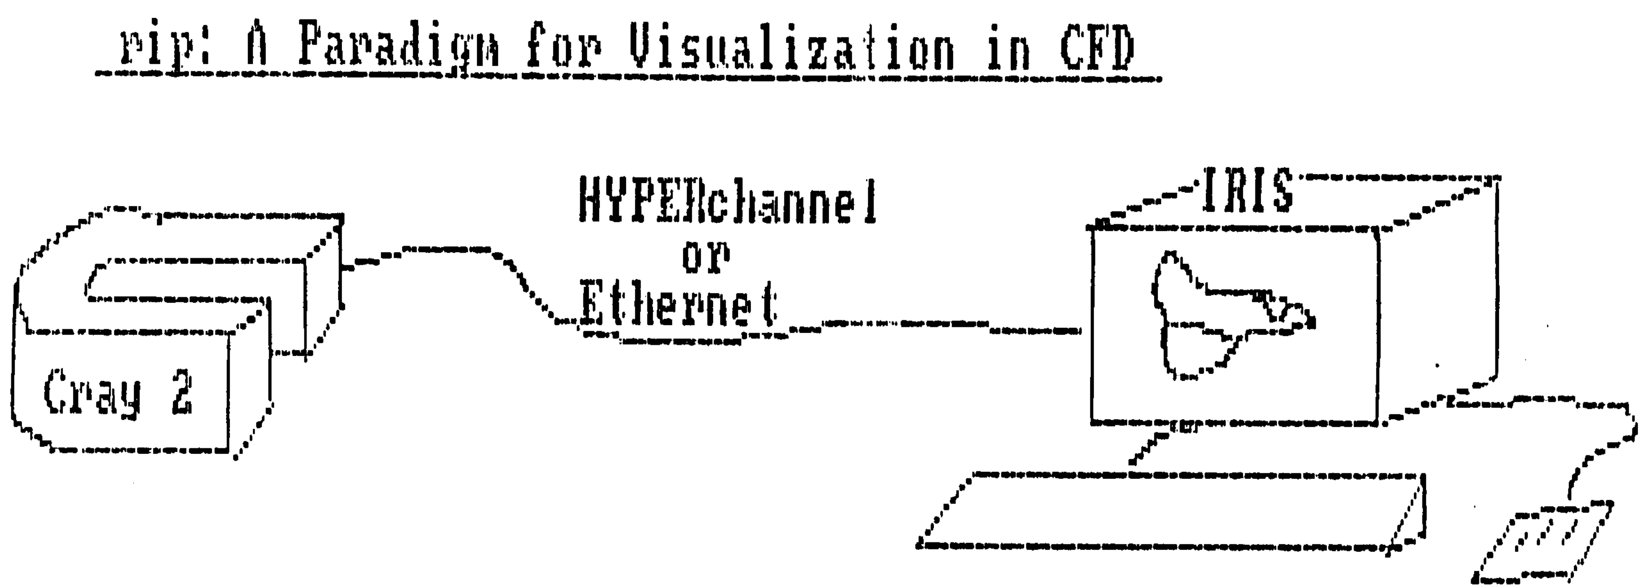

Fig. 3. An example of distributed scientific graphics.